arXiv:1611.06168v1 [stat.OT] 18 Nov 2016

On p-values

Laurie Davies
Faculty of Mathematics
University of Duisburg-Essen, 45117 Essen, Germany
email:laurie.davies@uni-due.de

1. The use and abuse of p-values

All branches of knowledge which require the analysis of data make use of p-values. Unfortunately in many cases 'make use of' could be replaced by 'abuse', the many reports of widespread abuse are convincing. In response The American Statistician published a statement on p-values by the American Statistical Association together with supplementary material consisting of statements by several statisticians and philosophers ([Wasserstein and Lazar, 2016]).

The most detailed of the supplementary material is [Greenland et al., 2016]. The authors point out that there are many ways in which any usefulness of a p-value can be invalidated. One example is to perform several experiments and report only the one with the smallest p-value. Problems of this nature will not be discussed here. It will be assumed that the experiment is so to speak clean and the data are so to speak high quality.

2. Probability models and approximation

2.1. **Semantics.** There are two meaning of the word 'model' is statistics. The first meaning refers to a parametric family of distributions. Thus the normal model is the family of all normal distributions. This meaning of the word 'model' is common in much of statistics, in particular in Bayesian statistics where such models are the objects of study.

The second meaning is that of a single probability measure. In this sense of the word the $N(0,1)$ probability measure is one model, the $N(0,2)$ probability measure is another model. This is the sense in which the word will be used in this paper. Models in this sense are the atoms so to speak of probability theory and hence the basic objects of stochastic modelling. The meaning of the word 'model' in the first sense is a parametric family of models in the second sense.

2.2. **Approximate models.** The authors of [Greenland et al., 2016] state

(A) .. the distance between the data and the model prediction is measured using a *test statistic* ..

and

(B) In logical terms, the p-value tests *all* the assumptions about how the data were generated (the entire model) ...

Although it is never precisely stated it seems that the word 'model' in the above quotations is meant in the first sense, a parametric family of distributions. Whatever the meaning of the word 'model' the meaning of the quotations taken together is clear. The distance between the data and the model is based on a test statistic and the corresponding p-value measures this distance in a particular manner. The phrase 'In logical terms' in the quotation (B) suggests that in practice this is not so. Indeed in practice the parametric model is accepted and the p-value is based on a particular hypothesis $H_0 : \mu = \mu_0$ using a statistic especially designed for testing this null hypothesis, for example a t-test. Such a single statistic cannot possibly test '*all* the assumptions about how the data were generated (the entire model)'.

A similar attitude is taken in [Birnbaum, 1962]: consideration is restricted to

(C) models whose adequacy is postulated and not in question ... the adequacy of any such model is typically supported, more or less adequately, by a complex informal synthesis of previous experimental evidence of various kinds and theoretical considerations concerning both subject-matter and experimental techniques.

In contrast to the word 'adequacy' being applied to a family of probability measures it will here be applied to individual probability measures. Thus the $N(0, 1)$ distribution may or may not be adequate for a given data set. The only sense I can make of applying the word 'adequate' to a parametric family of probability measures is that there are values of the parameter for which the individual distributions specified by these parameter values are consistent with the data. In general this will be a strict subset of the parameter space: it is difficult to imagine a data set for which the $N(0, 1)$ model and the $N(100, 10^{-6})$ model are both adequate or consistent with the data.

The two different meanings of the word 'model' are not just a question of notation or definition. They reflect two different approaches to statistics. This may be seen in [Birnbaum, 1962] where a parametric family of probability measures has

to be adequate without specifying the adequacy of any individual measure. This is necessary as the Likelihood Principle requires the proportionality of two different densities for all values of the parameter and not just for the adequate ones. A similar problem occurs when testing hypotheses. The parametric model is declared adequate without specifying the adequate values of the parameter. A hypothesis $H_0 : \mu = \mu_0$ is then tested to see whether $\mu = \mu_0$ is consistent with the data. It only makes sense to do this if the adequate parameter values have not been specified when declaring the whole family to be adequate as otherwise the test would be superfluous.

More generally a common approach is two perform a statistical analysis in two stages. In the first stage one or several parametric models will be investigated for adequacy, for example by using a goodness-of-fit test. Once an adequate model has been found it is made the basis of the second stage where it is treated as if it were true. Treating it as true means among other things ignoring the first stage. If indeed the model is now treated as if true then how we arrived at this truth is irrelevant. The following quotation is taken from the Chapter 5 of Huber [Huber, 2011] entitled 'Approximate Models':

(D)

> In the opposite case, if a goodness-of-fit test does not reject a model, statisticians have become accustomed to acting as if it were true. Of course this is logically inadmissible, even more so if with McCullagh and Nelder one believes that all models are wrong *a priori*.
> Moreover, treating a model that has not been rejected as correct can be misleading and dangerous. Perhaps this is the main lesson we have learned from robustness.

In [Davies, 2014] models are consistently treated as approximations. The basic idea is that a model P is an adequate approximation to a data set $\boldsymbol{x}_n = (x_1, \ldots, x_n)$ if typical data $\boldsymbol{X}_n(P)$ generated under P look like $\boldsymbol{x}_n$. Data are generated under single probability distributions P and not under a family of such, that is, a model in the first sense of the word. This is the reason why single probability distributions are the basic objects of study and not families of distributions.

The definition of 'look like' will depend on the nature of the data being analysed and the model. As an example suppose that the model is that of i.i.d. $N(0, 1)$ random variables. Then 'look like' can be based on the mean, the variance, the extreme values and the distance of the empirical distribution function from that of the standard $N(0, 1)$ distribution function. This will be done explicitly in Section 2.4

below. It is worth noting that the concept of adequacy is defined in terms of several statistics and not just one. This is in contrast to the quotation at the beginning of Section 3.2 where it is based on a single statistic.

The approach described in [Davies, 2014] can be read as an attempt to replace a two stage methodology, EDA followed by formal inference, by a single stage methodology whereby all tests become misspecification tests from without, or from a distance. It is an instance of 'distanced rationality' due to D. W. Müller (see [Müller, 1974]). Here my translation

(E)

> ... distanced rationality. By this we mean an attitude to the given, which is not governed by any possible or imputed immanent laws but which confronts it with draft constructs of the mind in the form of
>
> models, hypotheses, working hypotheses, definitions, conclusions, alternatives, analogies, so to speak from a distance, in the manner of partial, provisional, approximate knowledge.

2.3. **'Adequate' parametric families.** Although the quotation C does not make it explicit it is clear from [Birnbaum, 1962] that Birnbaum is referring to families of parametric models. Thus the Poisson family may be declared adequate without specifying any particular value of λ which is consistent with the data. As an example suppose the parametric family is the Poisson family and that the chi-squared goodness-of-fit is used to test adequacy. The test is typically based on some variant of the test statistic

$$\sum_{j=0}^{k} \frac{(\hat{p}_j - p_j(\bar{x}_n))^2}{p_j(\bar{x}_n)} \tag{1}$$

where the $\hat{p}_j$ are the empirical frequencies, $\bar{x}_n$ is the mean of the data and $p_j(\lambda) = \lambda^j \exp(-\lambda)/j!$. If the value of the test statistic (1) lies below a certain level then the Poisson model is declared adequate. Note that (1) does not specify any individual parameter values.

Given adequacy in this sense the whole parametric family is then transported to the second stage of formal inference in spite of the fact that the overwhelming majority of individual models will not be consistent with the data. For Birnbaum's argument to work this is essential: 'two likelihood functions, $f(x, \theta)$ and $g(y, \theta)$ are called the same if they are proportional, that is if there exists a positive constant c such that $f(x, \theta) = g(y, \theta)$ for all θ'.

In the second sense of the word 'model', an individual probability distribution, the goodness-of-fit procedure takes on a different form. In the concrete case of the Poisson family a given Poisson distribution say $\mathcal{P}_\lambda$ with $\lambda = 2$ can be tested for adequacy using

$$\sum_{j=0}^{k} \frac{(\hat{p}_j - p_j(2))^2}{p_j(2)} \,. \tag{2}$$

The set of λ values for which the test statistic lies below a critical level specifies those λ values, if any, which are consistent with the data. This will not be the set of all possible values.

If one interprets the concept of adequacy for models in the first sense using the second sense it can only mean that there are some parameter values θ for which the single model P_θ is consistent with the data. The likelihood principle is based on not specifying which values of θ these are.

2.4. **Approximation regions: an example.** A model P is an adequate approximation to data $\boldsymbol{x}_n$ if typical data sets generated under P look like $\boldsymbol{x}_n$. To make this susceptible to mathematical analysis the term 'look like' must be expressible in numerical quantities. This may not always be possible or easy. An animal may be easily recognizable as a dog but it it not easy to give this a mathematical expression. If a model is required which gives data sets looking like the daily returns of the Standard and Poor's 500 index it is not clear how 'look like' can be defined. In the following it will be assumed that 'look like' has a precise mathematical expression.

The following is taken from [Davies, 2014]. Given a probability measure P a sample of size n generated under P will be denoted by $\boldsymbol{X}_n(P) = (X_1(P), \ldots, X_n(P))$. Given further a family $\mathcal{P}$ of probability measures and a number $\alpha, 0 < \alpha \leq 1$, an α approximation region for the data $\boldsymbol{x}_n$ is defined by

$$\mathcal{A}(\boldsymbol{x}_n, \alpha, \mathcal{P}) = \{P \in \mathcal{P} : \boldsymbol{x}_n \in E_n(P)\} \tag{3}$$

where for each $P \in \mathcal{P}$ $E_n(P)$ denotes a subset of $\mathbb{R}^n$ such that

$$\boldsymbol{P}(\boldsymbol{X}_n(P) \in E_n(P)) = \alpha \,. \tag{4}$$

The choice of the $E_n(P)$ depends on the situation and has in general to be augmented by some form of regularization, for example: minimum Fisher models, number of local extremes, convexity constraints. These and further examples are to be found in [Davies, 2014].

The definition (3) makes no assumption that the data $\boldsymbol{x}_n$ were generated under some model $P_0 \in \mathcal{P}$. The interpretation is that $\mathcal{A}(\boldsymbol{x}_n, \alpha, \mathcal{P})$ specifies those models P for which $\boldsymbol{x}_n$ 'looks like' a 'typical sample' $\boldsymbol{X}_n(P)$ generated under P: typical samples $\boldsymbol{X}_n(P)$ lie in $E_n(P)$ so that points $\boldsymbol{x}_n \in E_n(P)$ look like typical samples $\boldsymbol{X}_n(P)$.

As an example suppose $\mathcal{P}$ is the family of normal distributions $\mathcal{N} = \{(\mu, \sigma) : N(\mu, \sigma^2)\}$. An approximation region can be based on the mean, the variance, outliers and the distance of the empirical measure to the model $N(\mu, \sigma^2)$ as measured by the Kuiper metric. More precisely put $\boldsymbol{y}_n = (\boldsymbol{x}_n - \mu)/\sigma$ and

$$(5) \qquad \begin{cases} T_1(\boldsymbol{y}_n) = \sqrt{n}\,|\mathrm{mean}(\boldsymbol{y}_n)|, & T_2(\boldsymbol{y}_n) = \sum_{i=1}^n y_i^2, \\ T_3(\boldsymbol{y}_n) = \max_i |y_i|, & T_4(\boldsymbol{y}_n) = d_{\mathrm{ku}}(\mathbb{P}(\boldsymbol{y}_n), N(0,1)), \end{cases}$$

where $\mathbb{P}(\boldsymbol{y}_n)$ is the empirical measure based on $\boldsymbol{y}_n$. Given $\tilde{\alpha}$ one can determine quantiles $q_1(\tilde{\alpha}), q_{21}(\tilde{\alpha}), q_{22}(\tilde{\alpha}), q_3(\tilde{\alpha}), q_4(\tilde{\alpha})$ such that

$$(6) \qquad \begin{cases} \boldsymbol{P}(T_1(\boldsymbol{Y}_n) \le q_1(\tilde{\alpha})) = \tilde{\alpha}, & \boldsymbol{P}(q_{21}(\tilde{\alpha}) \le T_2(\boldsymbol{Y}_n) \le q_{22}(\tilde{\alpha})) = \tilde{\alpha}, \\ \boldsymbol{P}(T_3(\boldsymbol{Y}_n) \le q_3(\tilde{\alpha})) = \tilde{\alpha}, & \boldsymbol{P}(T_4(\boldsymbol{Y}_n) \le q_4(\tilde{\alpha})) = \tilde{\alpha}. \end{cases}$$

where $\boldsymbol{Y}_n$ are i.i.d. $N(0,1)$. The approximation region is then defined by

$$(7) \quad \begin{aligned} \mathcal{A}(\boldsymbol{x}_n, \alpha, \mathbb{R} \times \mathbb{R}_+) \quad = \quad & \{(\mu, \sigma) : T_1(\boldsymbol{y}_n) \le q_1(\tilde{\alpha}), q_{21}(\tilde{\alpha}) \le T_2(\boldsymbol{y}_n) \le q_{22}(\tilde{\alpha}), \\ & T_3(\boldsymbol{y}_n) \le q_3(\tilde{\alpha}), T_4(\boldsymbol{y}_n) \le q_4(\tilde{\alpha}),\ \boldsymbol{y}_n = (\boldsymbol{x}_n - \mu)/\sigma\} \end{aligned}$$

where $\tilde{\alpha}$ is adjusted so that the region is indeed an α-approximation region. A reasonable starting value for $\tilde{\alpha}$ is $(3 + \alpha)/4$. This will lead to an effective value $\alpha^* > \alpha$ of α which can be determined by simulations. A better approximation can now be obtained by putting $\tilde{\alpha} = (3 + 2\alpha - \alpha^*)/4$. For a normal sample of size $n = 50$ and $\alpha = 0.9$ this leads to $\tilde{\alpha} \approx 0.97$ compared with the starting value of 0.975.

The following data give the quantity of copper in milligrams per litre in a sample of drinking water ([Davies, 2014]):

$$(8) \qquad \begin{aligned} & 2.16, 2.21, 2.15, 2.05, 2.06, 2.04, 1.90, 2.03, 2.06, 2.02, 2.06, 1.92, 2.08, \\ & 2.05, 1.88, 1.99, 2.01, 1.86, 1.70, 1.88, 1.99, 1.93, 2.20, 2.02, 1.92, 2.13, 2.13. \end{aligned}$$

The 0.9 approximation region $\mathcal{A}(\boldsymbol{x}_n, 0.9, \mathbb{R} \times \mathbb{R}_+)$ this data set is shown in Figure 1.

An approximation region for μ alone can be obtained by projecting $\mathcal{A}(\boldsymbol{x}_n, \alpha, \mathcal{N})$ onto the μ-axis:

$$(9) \qquad \mathcal{A}(\boldsymbol{x}_n, \alpha, \mathbb{R}) = \{\mu : \text{there exists some } \sigma \text{ s.t. } (\mu, \sigma) \in \mathcal{A}(\boldsymbol{x}_n, \alpha, \mathbb{R} \times \mathbb{R}_+)\}$$

This is equivalent to projecting the approximation region of Figure 1 onto the x-axis. The result is the interval $[1.945, 2.087]$. The standard 0.9 confidence interval for μ based on the t-statistic is the smaller interval $[1.978, 2.054]$. If the data really are normally distributed then the standard confidence interval for μ will be smaller than the corresponding approximation interval. If the data are not normally distributed then the approximation interval can be smaller, indeed much smaller than the confidence interval. This will be discussed in Section 2.6 below.

In (7) the same $\tilde{\alpha}$ is used for all four functionals. There is no need for this. If for example the Kuiper distance is not regarded as important as the other features it can be given less weight in terms of a higher value of $\tilde{\alpha}$.

2.5. **Multiple p-values.** The approximation region (7) is based on the four statistics $T_i, i = 1, \ldots, 4$. For each parameter pair (μ, σ) each of the statistics $T_i, i = 1, 3, 4$ comes with a p-value

$$p_i(\mu, \sigma) = 1 - \boldsymbol{P}(T_i(\boldsymbol{Y}_n) \leq T_i(\boldsymbol{y}_n)), \tag{10}$$

the statistic T_2 comes with the p-value

$$p_2(\mu, \sigma) = 2\min(\boldsymbol{P}(T_2(\boldsymbol{Y}_n) \leq T_2(\boldsymbol{y}_n)), 1 - \boldsymbol{P}(T_2(\boldsymbol{Y}_n) \leq T_2(\boldsymbol{y}_n))) \tag{11}$$

where the $\boldsymbol{Y}_n$ are i.i.d. $N(0, 1)$ and $\boldsymbol{y}_n = (\boldsymbol{x}_n - \mu)/\sigma)$. Thus each parameter pair (μ, σ) comes with four p-values attached $p_i(\mu, \sigma), i = 1, \ldots 4$. It belongs to the approximation region if and only if

$$p(\mu, \sigma) = \min(p_i(\mu, \sigma), i = 1, \ldots, 4) \geq 1 - \alpha^*. \tag{12}$$

As an example the pair $(2.008, 0.110)$ in the approximation region of Figure 1 has the p_i-values $(0.720, 0.683, 0.123, 0.967)$.

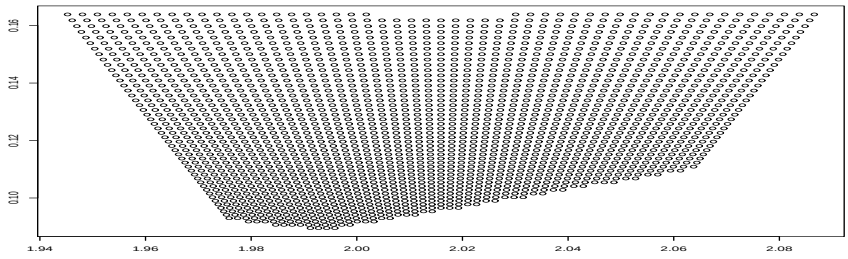

FIGURE 1. The approximation region $\mathcal{A}(\boldsymbol{x}_n, 0.9, \mathbb{R} \times \mathbb{R}_+)$ for the data (8).

The multiple p-values associated with each parameter value stand in contrast to the usual definition of a p-value which uses only one statistic (see the quotation A).

2.6. **Approximation and confidence regions.** At first sight the approximation region (3) can be interpreted as a confidence region. If the data $\boldsymbol{x}_n$ were indeed generated under some model $P_0 \in \mathcal{P}$ then because of (4) we have

$$\boldsymbol{P}(P_0 \in \mathcal{A}(\boldsymbol{X}_n(P_0), \alpha, \mathcal{P})) = \alpha \text{ for all } P_0 \in \mathcal{P}\,. \tag{13}$$

Such an interpretation however causes difficulties. Consider the family $\mathcal{P} = \{N(\mu, \sigma^2) : (\mu, \sigma) \in \mathbb{R} \times \mathbb{R}_+\}$. A standard confidence region for the 'true' value μ_0 of μ is based on the assumption that there is indeed a 'true' value μ_0 of μ. That is the data were generated under $N(\mu_0, \sigma^2)$ for some σ. This assumption is not checked in the formal inference phase and consequently a confidence region for μ_0 is never empty. The interpretation is that it is a measure precision with which μ_0 can be determined.

The approximation region (9) on the other hand is not based on the assumption that the data were indeed generated as i.i.d. $N(\mu, \sigma^2)$ for some (μ, σ). It specifies those μ-values if any for which $N(\mu, \sigma^2)$ is an adequate approximation to the data for some σ. Thus if the adequacy region (9) is small this simply means that there are few values of μ for which $N(\mu, \sigma^2)$ is an adequate approximation to the data for some σ. It is not a measure of precision. If one imagines the data gradually becoming less and less normal then the region (7) will become smaller and eventually will be the empty set. One way of doing this is to gradually increase one value of the sample until this value becomes incompatible with the feature T_3 of (5). As an example Figure 2 gives the 0.9 approximation region for the copper data of (8) but with the smallest observation of 1.7 being replaced by 1.5 If the 1.7 is replaced by 1.267 the

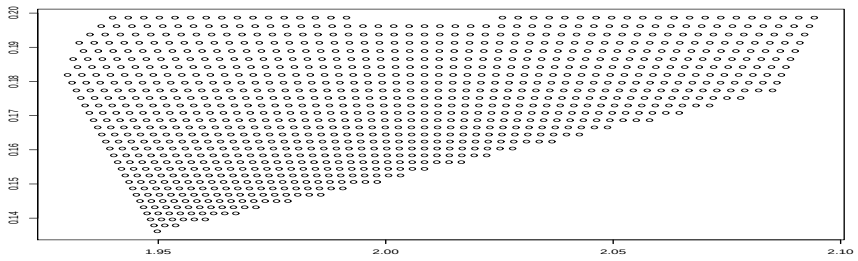

FIGURE 2. The approximation region $\mathcal{A}(\boldsymbol{x}_n, 0.9, \mathbb{R} \times \mathbb{R}_+)$ for the data (8) but with 1.7 replaced by 1.5.

approximation region as calculate has exactly one point $(1.9812, 0.2161)$.

Interpreting (9) as a confidence region leads to complications. As the data become less and less like Gaussian data the region becomes smaller and smaller which is interpreted as an increase in precision. Thus on this interpretation replacing 1.7 by 1.267 in (8) leads to exact values for (μ, σ) namely $(1.9812, 0.2161)$. When the region becomes empty this is as if one goes from infinite precision to no information at all. A discussion can be found in

(*) http://andrewgelman.com/2011/08/25/ .

From the point of view of approximation there is no problem of interpretation. The set of adequate parameter values becomes smaller and smaller and eventually becomes the empty set, that is, there are no adequate parameter values at all.

2.7. **An empty approximation region.** Consider the approximation region (7) with $\tilde{\alpha} = 0.975$ corresponding to $\alpha \approx 0.92$. Simulations show that for normal samples of size 50 the approximation is empty in about 0.7% of the cases. This value is based on 5000 simulations. It is much smaller than the 8% of the cases where the approximation region does not contain the (μ, σ) pair used to generate the data.

If the approximation region is empty, that is, the family $\mathcal{P}$ contains no model which is an adequate approximation, there may well be an interest in quantifying just how poor the approximation is. One way of doing this is to determine the smallest value of α, say α^* such that the approximation region is non-empty. The corresponding p-value is defined as $p^* = 1 - \alpha^*$ which is a measure of the goodness of the approximation: the smaller the p-value the worse the approximation.

For the approximation region (7) it is always possible to calculate p as the quantiles q can be calculated. If the quantiles were obtained by simulation and the approximation is poor then it may not be possible to calculate the p-values. An alternative is suggested in [Lindsay and Liu, 2009]. It is based on the idea that it is easier for there to be an adequate approximation if the sample size is small. Samples $\boldsymbol{x}_m^*$ of size m are drawn for the original sample $\boldsymbol{x}_n$ and the approximation region calculated. The measure of the degree of approximation is the largest value of m for which the approximation region is not empty in 50% of the cases. An example is given in Chapter 3.8 of [Davies, 2014] for a sample of size $n = 189$. The family of models considered was the family of discretized gamma models and the concept of adequacy was based on the total variation metric The fit was so poor that even for $\alpha = 0.999999$ there was no adequate approximation. The size of the smallest

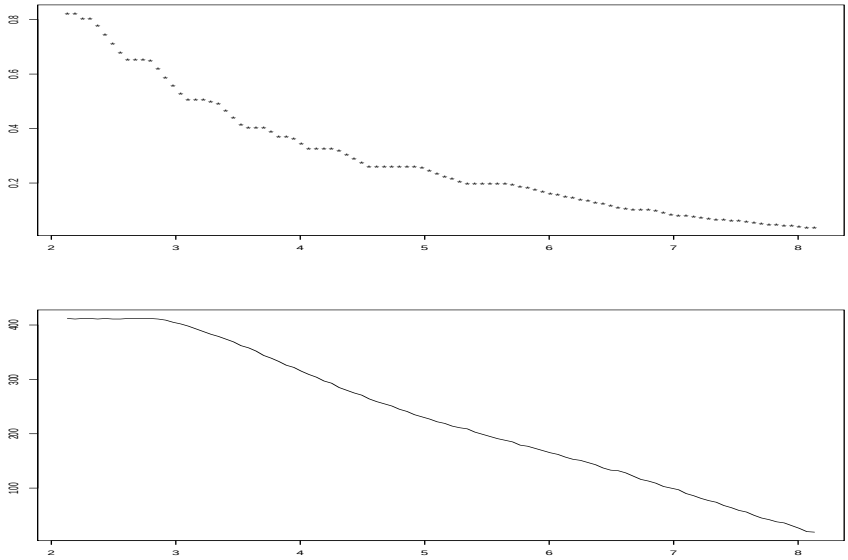

FIGURE 3. Upper panel: the p-values (12) plotted against the of the largest observation for a normal sample of size $n = 27$. Lower panel: the number of points in the approximation region also plotted against the size of the largest observation.

random subsamples for which there was an adequate approximation in 50% of the cases was approximately 40.

2.8. **A non-empty approximation region.** If the approximation region is defined by statistics as in (7) then for each model P the p-values for each of inequalities is calculated and the minimum value taken. The maximum of these values taken over the approximation region is then a measure of the degree of approximation. Again, the smaller this value the poorer the approximation. The statistician can base the decision on whether to use the family of models $\mathcal{P}$ at least in part on the maximum p-value. A cut-off point 0.2 implies that there are adequate models where the p-values of the functionals involved all exceed 0.2.

Figure 3 shows an example of this. A $N(0, 1)$ sample of size $n = 50$ was generated and the largest observation 2.130 gradually increased in steps of 0.06875. The upper panel shows the p-values as a function of the size of this observation, the lower panel the the number of points in the approximation region evaluated over a grid of parameter values. It can be interpreted as a proxy for the area of the region.

The quantiles or the p-values of the p-values can be obtained from simulations. For a normal sample of size $n = 28$ the 0.001, 0.01, 0.05 and 0.1 quantiles are 0.0713, 0.210, 0.324 and 0.393 respectively. These values are based on 10000 simulations.

The p-value based on the normal family of models is 0.406 with a p-value (p-value of the p-value) of about 0.1. If the smallest observation 1.70 is removed the p-value becomes 0.835 which corresponds to a p-value of 0.96. If the smallest value is set to 1.4 the p-value of the p-values is about 0.001. This raises the questions to how bad an approximation can be whilst still basing an analysis on the family of distributions.

The following quotation from [Huber, 2011] is relevant in this context. It immediately precedes the quotation D

(F)

> If a goodness-of-fit tests rejects a model, we are left with many alternative actions. Perhaps we do nothing (and continue to live with a model certified to be inaccurate). Perhaps we tinker with a model by adding or adjusting a few more features (and thereby destroy its conceptual simplicity). Or we switch to an entirely different model, maybe one based on a different theory, or maybe in the absence of such a theory, to a purely phenomenological one. In the opposite case, if a goodness-of-fit test does not reject a ...

2.9. **p-values and hypotheses.** Consider the Gaussian family of models and the null hypothesis $H_0 : \mu \geq \mu_0$. The p-value defined by the t-statistic

$$\text{pt}\left(\sqrt{n}\,\frac{\text{mean}(\boldsymbol{x}_n) - \mu_0}{\text{sd}(\boldsymbol{x}_n)}, n-1\right) \tag{14}$$

is often used as a measure as to the extent that μ_0 is compatible with the data. This definition is not acceptable from the point of view of approximation as it does not specify any value for σ.

As an example consider the copper data (8) Suppose the legal limit is 2.1 milligrams per litre and we wish to test the hypothesis that this is exceeded. The Gaussian family of models will be used so that with the usual identification of the amount of copper in the water with μ the null hypothesis becomes

$$H_0 : \mu \geq 2.1. \tag{15}$$

The p-value as defined by (14) is 0.000436.

An equivalent definition of a standard p-value is the following. Given the p-value p^* put $\alpha^* = 1 - p^*$. Then α^* is the smallest value of α such that the confidence

region for μ contains μ_0. This can be used to define a p-value using the idea of an approximation region. This p-value is defined as $p^* = 1-\alpha^*$ where α^* is the smallest value of α such that the approximation region contains a point (μ_0, σ) for some σ. This is similar to the definition of a p-value for an empty approximation region given in Section 2.7 . If this is done for the water data (8) using the approximation region (7) the resulting p-value is 0.045.

Replace now the smallest value 1.7 by 0.7. The p-value of (14) is now 0.015. At first sight this may seem surprising as the value 0.7 is less consistent with (15) than is 1.7. The reason is that the standard deviation is now 0.274 as against 0.116. The p-value based on the approximation region is 0.00018. The value of σ is 0.310. The reason is that the value 0.7 is essentially an outlier. This is picked up by the statistics T_3 and T_4 but not by the t-statistic. See the second Huber quotation (F). The outlyingness of 0.7 should have been detected in the EDA phase before moving on to the formal inference phase. This raises the question of how to react to the outlier.

3. p-VALUES AND FUNCTIONALS

The purpose of the copper measurements (8) it to give a point estimate of the amount of copper in the sample of drinking water combined with an interval of reasonable values. The mean and a confidence interval using a normal model give a reasonable solution for this particular data sets, but there are problems.

One immediate question is why the Gaussian family and not the Laplace (double exponential) family? This question draws attention to the fact that the location-scale problem is ill-posed when density based methods, maximum likelihood or Bayes, are used. Some form of regularization is required. What are required are 'bland' or 'hornless' models (see Section 2, *B is for Blandness*, of [Tukey, 1993] and Chapter 1.3.6 of [Davies, 2014]). In the location-scale situation one possible form of regularization is to use minimum Fisher information models such as the Gaussian.

Another problem is to relate the parameters of the model to the real world. As the purpose of the copper data is to estimate the amount of copper in the water, simply estimating (in another sense of the word estimate) the parameters of a parametric model does not solve the problem. The parameters must be connected to the real world. For the Gaussian family this is not a problem as the canonical connection is to identify the location parameter μ with the actual amount of copper in the water. However this fails for the log-normal distribution, another minimum

Fisher distribution. One can still associate the actual amount of copper with the mean but also with the median. This gives two different identifications for the same model.

The final problem is that of outliers. They are common in interlaboratory tests and any method of analysis must be able to deal with them. Neither the Gaussian, Laplace or log-normal achieve this.

The path taken in Chapter 5 of [Davies, 2014] is to use M-functionals (see Chapters 4 and 5 of [Huber and Ronchetti, 2009]). Given ψ- and χ-functions ψ and χ respectively and a probability measure P over $\mathbb{R}$ the M-functional T_M is defined by $T_M(P) = (T_L(P), T_S(P)$ where $T_L(P)$ and $T_S(P)$ solve

$$(16) \qquad \begin{cases} \int \psi\left(\frac{x-T_L(P))}{T_S(P)}\right) dP(x) &= 0 \\ \int \chi\left(\frac{x-T_L(P))}{T_S(P)}\right) dP(x) &= 0 \end{cases}$$

The functions ψ and χ can be so chosen so that (i) (16) has a unique solution for all P with a largest atom of less than 0.5 and (ii) the functional $T_M(P)$ is locally uniformly Fréchet differentiable in a Kolmogorov neighbourhood of P see([Davies, 1998] and page 54 of [Hampel et al., 1986]). This gives stability of analysis with respect to P. The functions used here are

$$(17) \qquad \psi(u,c) = \frac{\exp(u/c)-1}{\exp(u/c)+1}, \quad \chi(u) = \frac{u^4-1}{u^4+1}$$

where c is a tuning constant set here to 5.

The connection with reality is achieved by identifying the amount of copper with $T_L(P)$ for any adequate model P. the only form of adequacy required is that the model P is in a small Kolmogorov neightbourhood of the empirical distribution $\mathbb{P}_n$ of the data. This still leaves open the choice of T_M. This will be discussed at the end of the section.

Let $\boldsymbol{X}_n(P)$ denote a sample of size n of i.i.d. random variable with distribution P, and by $\mathrm{q}_\psi(\cdot, n, P)$ the quantiles of

$$\frac{1}{\sqrt{n}}\left|\sum_{i=1}^{n} \psi\left(\frac{X_i(P)-T_L(P)}{T_S(P)}\right)\right|$$

with the corresponding definition of $\mathrm{q}_\chi(\cdot, n, P)$. The an α-approximation region for the functional T_M is defined by

$$\begin{aligned}(18)\quad \mathcal{A}(\boldsymbol{x}_n, \alpha, T_M) \;=\; &\Big\{(T_L(P), T_S(P)) : d_{ko}(\mathbb{P}_n, P) < \mathrm{qdk}(\tilde{\alpha}, n),\\ &\sqrt{n}\left|\int \psi\left(\frac{x - T_L(P)}{T_S(P)}\right) d\mathbb{P}_n(x)\right| \le \mathrm{q}_\psi(\tilde{\alpha}, n, P),\\ &\sqrt{n}\left|\int \chi\left(\frac{x - T_L(P)}{T_S(P)}\right) d\mathbb{P}_n(x)\right| \le \mathrm{q}_\chi(\tilde{\alpha}, n, P)\Big\}\end{aligned}$$

where $\tilde{\alpha} = (2 + \alpha)/3$, $\mathbb{P}_n$ denotes the empirical distribution of the data $\boldsymbol{x}_n$, d_{ko} the Kolmogorov metric and $\mathrm{qdk}(\cdot, n)$ its quantile function. The choice of $\tilde{\alpha}$ corresponds to spending $(1 - \alpha)/3$ on each of the three features in the definition of the approximation region.

As

$$\boldsymbol{E}\left(\psi\left(\frac{X_i(P) - T_L(P)}{T_S(P)}\right)\right) = 0$$

and

$$\boldsymbol{E}\left(\psi\left(\frac{X_i(P) - T_L(P)}{T_S(P)}\right)^2\right) = \int \psi\left(\frac{x - T_L(P)}{T_S(P)}\right)^2 dP(x)$$

it follows from the central limit theorem that

$$\frac{1}{\sqrt{n}}\sum_{i=1}^{n} \psi\left(\frac{X_i(P) - T_L(P)}{T_S(P)}\right) \asymp N\left(0, \int \psi\left(\frac{x - T_L(P)}{T_S(P)}\right)^2 dP(x)\right).$$

with the same result for χ. Thus asymptotically

$$(19)\qquad \mathrm{q}_\psi(\tilde{\alpha}, n, P) \approx \mathrm{qnorm}(\tilde{\alpha})\sqrt{\int \psi\left(\frac{x - T_L(P)}{T_S(P)}\right)^2 dP(x)}$$

with the same result for χ. As the random variables are bounded the normal approximation is good for small values of n.

The requirement $d_{ko}(\mathbb{P}_n, P) < \mathrm{qdk}(\tilde{\alpha}, n)$ forces P into a $O(1/\sqrt{n})$ Kolmogorov neighbourhood of $\mathbb{P}_n$. This together with the locally uniform Fréchet differentiability implies $\mathrm{q}_\psi(\tilde{\alpha}, n, P) \approx \mathrm{q}_\psi(\tilde{\alpha}, n, \mathbb{P}_n)$ (see pages 107-108 of[Davies, 2014]) and together with (19) it leads to the approximate approximation region

$$\begin{aligned}(20)\qquad \tilde{\mathcal{A}}(\boldsymbol{x}_n, \alpha, T_M) = &\Big\{(T_L(P), T_S(P)) : d_{ko}(\mathbb{P}_n, P) < \mathrm{qdk}(\tilde{\alpha}, n)\\ &\sqrt{n}\left|\int \psi\left(\frac{x - T_L(P)}{T_S(P)}\right) d\mathbb{P}_n(x)\right| \le \mathrm{qnorm}(\tilde{\alpha})\sqrt{\boldsymbol{V}_\psi(\mathbb{P}_n)},\\ &\sqrt{n}\left|\int \chi\left(\frac{x - T_L(P)}{T_S(P)}\right) d\mathbb{P}_n(x)\right| \le \mathrm{qnorm}(\tilde{\alpha})\sqrt{\boldsymbol{V}_\chi(\mathbb{P}_n)}\end{aligned}$$

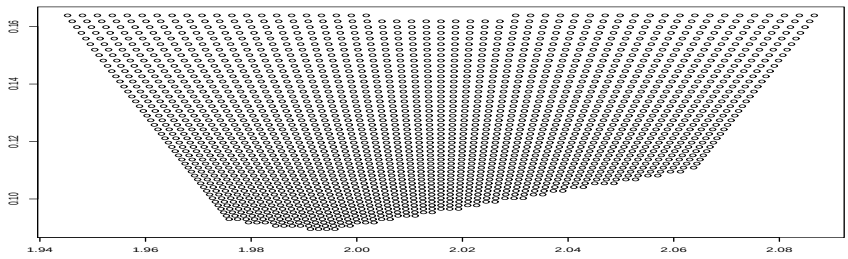

FIGURE 4. The 0.9 approximation region of the location and scale functionals (T_L, T_S) for the copper data using the psi- and chi-functions of (17) with $c = 5$.

where

$$\boldsymbol{V}_\psi(\mathbb{P}_n) = \frac{1}{n}\sum_{i=1}^{n} \psi\left(\frac{x_i - T_L(P)}{T_S(P)}\right)^2 .$$

This approximation to (18) can be calculated over a grid of values. It is shown in Figure 4 for the copper data with $\alpha = 0.9$. It may be compared with the approximation region based on the Gaussian distribution as shown in Figure 1. The approximation region for $T_L(P)$ is obtained by projecting the approximation region onto the x-axis. For the copper data with $\alpha = 0.9$ it is $[1.964, 2.067]$ compared with the standard 0.9-confidence interval $[1.978, 2.054]$ based on the t-statistic.

The approximation region (4) remains unchanged if the smallest observation 1.7 is replaced by zero. This is in sharp contrast to the approximation region based on the Gaussian family of models which is empty in this case. This one example of stability of analysis deriving from the use of T_M: small changes in the data, here a single data point, lead to only small changes in the result. It was pointed out above that the location-scale problem requires regularization. The use of the M-functional T_M is a regularization of the procedure not the models.

Hypothesis testing as in Section 2.9 can be done as follows. For the copper data the null hypothesis (15) is replaced by

$$H_0 : T_L(P) \geq 2.1.$$

The p-value is $p^* = 1 - \alpha^*$ where α^* is the smallest value of α such that (20) contains $(2.1, \sigma)$ for some σ. Its value is $p^* = 0.01$.

The M-functional used here is not the only one. There are many possible choices. Which one to use is an empirical question. A member of the committee which produced the German DIN standard ([DIN, 2003]) for analysing water, waste water

and sludge reported that in his experience the median was better than the mean but worse than the mean after the elimination of outliers. The final decision was to use Hampel's redescending ψ-function (Example 1 on page 150 of [Hampel et al., 1986]) which can be seen as a smooth version of the mean after eliminating outliers.

4. Approximation and prediction

4.1. **Prediction.** The concept of adequate approximation can be looked at in terms of prediction. Given a number α and based on a model P a prediction has to be made about a sample $\boldsymbol{x}_n$. That the prediction is based on P means that if the sample were generated under P, that is $\boldsymbol{x}_n = \boldsymbol{X}_n(P)$, then the prediction would be correct with probability α. In making the prediction is has to be decided which aspects of the data are regarded as important. In the definition of the approximation region (7) the important aspects are given by the statistics $T_i, i = 1, \ldots, 4$. With $P = N(\mu, \sigma^2)$ the corresponding prediction is that all the inequalities of (6) will hold with $\boldsymbol{y}_n = (\boldsymbol{x}_n - \mu)/\sigma$ replacing $\boldsymbol{Y}_n$. If the prediction is correct then the model $N(\mu, \sigma^2)$ is accepted as an adequate approximation to the data.

4.2. **Jeffreys on p-values.** The following is often cited as an argument against the use of p-values:

(G) .... gives the probability of departures, measured in a particular way, equal to or greater than the observed set, and the contribution from the actual value is nearly always negligible. What the use of P implies, therefore, is that a hypothesis that may be true may be rejected because it has not predicted observable results that have not occurred. This seems to be a remarkable procedure. On the face of it, the evidence might more reasonably be taken as evidence for the hypothesis, not against it.

(page 385 of [Jeffreys, 1961]).

Suppose the hypothesis is that the data follow the $N(0, 1)$ distribution. What observable results does this hypothesis predict? It seem pointless to predict a single value as such a prediction would be wrong with probability 1. The prediction must be a set $\mathcal{S}$ of values with the prediction being regarded as correct if the observable result x lies in $\mathcal{S}$. Putting $\mathcal{S} = \mathbb{R}$ results in the prediction being correct with probability one but this is somewhat vacuous. A non-vacuous prediction can be obtained by specifying a probability α and a set $\mathcal{S}(\alpha)$ such that the prediction is

correct with probability α, $\boldsymbol{P}(X \in \mathcal{S}(\alpha)) = \alpha$. It is worthy of note that the larger α the more vacuous the prediction so to speak. As a simple example put $\alpha = 0.95$ and $\mathcal{S}(\alpha) = (-1.96, 1.96)$ and suppose that $x = 3.121$ is observed. The p-value is $\boldsymbol{P}(|X| > 3.121) = 0.0018$ and for this to be a successful prediction would require $\alpha = 0.9982$ rather than the chosen $\alpha = 0.95$. We now interpret 'not predicted to occur' in the sense 'predicted not to occur' rather than in the sense 'forgetting to predict'. If it were agreed beforehand that a false prediction would lead to the null hypothesis to be rejected, then this is done because a value predicted not to occur, namely 3.121, did in fact occur. This seems an unremarkable procedure. How bad the prediction error is can be measured by the $\alpha = 0.9982$ required to make the prediction correct and which corresponds to a very weak prediction in that it would be correct in 99.8% of the times.

5. p-VALUES AND CHOICE OF COVARIATES IN STEPWISE REGRESSION

The following is based on [Davies, 2016a]. Given a data set of size n consisting of a dependent variable $\boldsymbol{y}(n)$ and $p(n)$ covariates $\boldsymbol{x}(n)$ the problem is to decide which if any of the covariates to include. The discussion below will be restricted to the case where $p(n)$ is chosen by stepwise regression but the idea can be extended to considering all subsets of the covariates as long as $p(n)$ is not too large, say $p(n) \leq 20$ (see [Davies, 2016b]).

It would seem that all procedures for choosing the covariates are based on the standard linear model

$$\boldsymbol{Y}(n) = \boldsymbol{X}(n)\boldsymbol{\beta}(n) + \boldsymbol{\varepsilon}(n). \tag{21}$$

The procedure to be described below is not based on this model. The basic idea is to compare the covariates $\boldsymbol{x}(n)$ with covariates which are simply standard Gaussian white noise. A covariate $\boldsymbol{x}_j$ is included only if it is significantly better than white noise.

Suppose that $p_0 \leq n-2$ with indices $\mathcal{S}_0$ have already been been included in the regression and that the sum of squared residuals is ss_0. Denote by ss_j the sum of squared residuals if the covariate $\boldsymbol{x}_j$ with $j \notin \mathcal{S}_0$ is included. The next candidate for inclusion is that covariate for which ss_j is smallest. Including this covariate leads to a sum of squared residuals

$$ss_{01} = \min_{j \notin \mathcal{S}_0} ss_j.$$

Replace all the covariates not in $\mathcal{S}_0$ in their entirety by standard Gaussian white noise. Let SS_j denote the sum of squared residuals if the random covariate corresponding to $\boldsymbol{x}_j$ is included. The inclusion of the best of the random covariates leads to a sum of squared residuals

$$SS_{01} = \min_{j \notin \mathcal{S}_0} SS_j.$$

The probability that the best random covariate is better than the best of the actual covariates is

$$\begin{aligned} \boldsymbol{P}(SS_{01} < ss_{01}) &= 1 - \boldsymbol{P}(SS_{01} \geq ss_{01}) = 1 - \boldsymbol{P}(\min_{j \notin \mathcal{S}_0} SS_j \geq ss_{01}) \\ &= 1 - \prod_{j \notin \mathcal{S}_0} \boldsymbol{P}(SS_j \geq ss_{01}) \end{aligned}$$

It has been shown by Lutz Dümbgen that

$$SS_j \stackrel{D}{=} ss_0(1 - B_{1/2,(n-p_0-1)/2}) \tag{22}$$

where $B_{a,b}$ denotes a beta random variable with parameters a and b and distribution function $\text{pbeta}(\cdot, a, b)$. From this it follows that

$$\boldsymbol{P}(SS_j \geq ss_{01}) = \text{pbeta}(1 - ss_{01}/ss_0, 1/2, (n - p_0 - 1)/2)$$

so that finally

$$\boldsymbol{P}(SS_{01} \leq ss_{01}) = 1 - \text{pbeta}(1 - ss_{01}/ss_0, 1/2, (n - p_0 - 1)/2)^{p(n)-p_0}. \tag{23}$$

This is the p-value for the inclusion of the next covariate. The simplest procedure is to specify $\alpha < 1$ and to continue the stepwise selection until the first p-value exceeds α. Those covariates up to but excluding this last one are the selected ones. The stopping rule is

$$ss_{01} > ss_0 \left(1 - \text{qbeta}((1-\alpha)^{1/(p(n)-p_0)}, 1/2, (n - p_0 - 1)/2)\right) \tag{24}$$

where $\text{qbeta}(\cdot, a, b)$ is the quantile function of the beta distribution with parameters a and b.

The procedure is conceptually and algorithmically simple. It requires no regularization parameter or cross-validation or an estimate of the error term in (21). It is invariant with respect to affine changes of unit of the covariates and equivariant with respect to a permutation of the covariates. It can be extended to non-linear parametric regression, robust regression and the Kullback-Leibler discrepancy where appropriate.

As an example we take the leukemia data ([Golub et al., 1999]

http://www-genome.wi.mit.edu/cancer/

which was analysed in [Dettling and Bühlmann, 2003]. These consist of data on $n = 72$ samples of tissue with with $p(n) = 3571$ covariates. The dependent variable $\boldsymbol{y}(n)$ is either 0 or 1 depending on whether the patient suffers from acute lymphoblastic leukemia or acute myeloid leukemia. The first five genes in order of inclusion with their associated p-values as defined by (23) are as follows:

(25)

gene number	1182	1219	2888	1946	2102
p-value	0.0000	8.57e-4	3.56e-3	2.54e-1	1.48e-1

According to this relevant genes are 1182, 1219 and 2888 and given these the remaining 3568 are no better than random noise. This applies to the gene 1946 but if a simple linear regression is performed using this gene alone its p-value in the linear regression is 7.75e-9. This is much smaller than the 0.254 in (25). The p-value (23) takes into account the stepwise nature of the procedure, in particular that gene 1946 is the best of the remaining genes once the genes 1182, 1219 and 2888 have been included. A simple linear regression does not take this into account.

The data were gathered in the hope of using the gene expression data to classify the patients. If the classification is based on genes 1182, 1219 and 2888. A simple linear regression results in one misclassification. In [Dettling and Bühlmann, 2003] the authors considered 42 different classification schemes. Only two of them resulted in a single misclassification. They used a 1-nearest-neighbour method based on 25 and 3571 genes. For this particular data set the procedure described above attains the same result and moreover specifies the relevant genes.

References

[DIN, 2003] (2003). DIN 38402-45:2003-09 German standard methods for the examination of water, waste water and sludge - General information (group A) - Part 45: Interlaboratory comparisons for proficiency testing of laboratories (A 45). Deutsches Institut für Normierung.

[Birnbaum, 1962] Birnbaum, A. (1962). On the foundations of statistical inference. *Journal of the American Statistical Association*, 57:269–326.

[Davies, 2014] Davies, L. (2014). *Data Analysis and Approximate Models*. Monographs on Statistics and Applied Probability 133. CRC Press.

[Davies, 2016a] Davies, L. (2016a). Stepwise choice of covariates in high dimensional regression. arXiv:1610.05131 [math.ST].

[Davies, 1995] Davies, P. L. (1995). Data features. *Statistica Neerlandica*, 49:185–245.

[Davies, 1998] Davies, P. L. (1998). On locally uniformly linearizable high breakdown location and scale functionals. *Annals of Statistics*, 26:1103–1125.

[Davies, 2008] Davies, P. L. (2008). Approximating data (with discussion). *Journal of the Korean Statistical Society*, 37:191–240.

[Davies, 2016b] Davies, P. L. (2016b). Functional choice and non-significance regions in regression. arXiv:1605.01936 [math.ST].

[Dettling and Bühlmann, 2003] Dettling, M. and Bühlmann, P. (2003). Boosting for tumor classification with gene expression data. *Bioinformatics*, 19(9):1061–1069.

[Golub et al., 1999] Golub, T., Slonim, D., P., T., Huard, C., Gaasenbeek, M., Mesirov, J., Coller, H., Loh, M., Downing, J., Caligiuri, M., Bloomfield, C., and Lander, E. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science*, 286(15):531–537.

[Greenland et al., 2016] Greenland, S., Senn, S. J., Rothman, K. J., Carlin, J. B., Poole, C., Goodman, S. N., and Altman, D. G. (2016). Statistical tests, *p*-values, confidence intervals, and power: A guide to misinterpretations. The American Statistician, Volume 70. Supplemental material to 'The ASA's Statement on p-Values: Context, Process and Purpose'.

[Hampel et al., 1986] Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., and Stahel, W. A. (1986). *Robust Statistics: The Approach Based on Influence Functions*. Wiley, New York.

[Huber, 2011] Huber, P. J. (2011). *Data Analysis*. Wiley, New Jersey.

[Huber and Ronchetti, 2009] Huber, P. J. and Ronchetti, E. M. (2009). *Robust Statistics*. Wiley, New Jersey, second edition.

[Jeffreys, 1961] Jeffreys, H. (1961). *Theory of Probability*. Oxford Classic Texts in the Physical Sciences. Oxford University press, third edition.

[Lindsay and Liu, 2009] Lindsay, B. and Liu, J. (2009). Model assessment tools for a model false world. *Statistical Science*, 24(3):303–318.

[Müller, 1974] Müller, D. W. (1974). Thesen zur Didaktik der Mathematik. *Math. phys. Semesterberichte, N.F.*, 21:164–169.

[Tukey, 1993] Tukey, J. W. (1993). Issues relevant to an honest account of data-based inference, partially in the light of Laurie Davies's paper. Princeton University, Princeton.

[Wasserstein and Lazar, 2016] Wasserstein, R. L. and Lazar, N. A. (2016). The asa's statement on p-values: Context, process, and purpose. *The American Statistician, Volume 70*, 70(2).